

Saliency Driven Image Manipulation

Roey Mechrez
Technion

roey@tx.technion.ac.il

Eli Shechtman
Adobe Research

elishe@adobe.com

Lihi Zelnik-Manor
Technion

lihi@ee.technion.ac.il

Abstract

Have you ever taken a picture only to find out that an unimportant background object ended up being overly salient? Or one of those team sports photos where your favorite player blends with the rest? Wouldn't it be nice if you could tweak these pictures just a little bit so that the distractor would be attenuated and your favorite player will stand-out among her peers? Manipulating images in order to control the saliency of objects is the goal of this paper. We propose an approach that considers the internal color and saliency properties of the image. It changes the saliency map via an optimization framework that relies on patch-based manipulation using only patches from within the same image to achieve realistic looking results. Applications include object enhancement, distractors attenuation and background decluttering. Comparing our method to previous ones shows significant improvement, both in the achieved saliency manipulation and in the realistic appearance of the resulting images.

1. Introduction

Saliency detection, the task of identifying the salient and non-salient regions of an image, has drawn considerable amount of research in recent years, e.g., [14, 18, 21, 30, 32]. Our interest is in manipulating an image in order to modify its corresponding saliency map. This task has been named before as *attention retargeting* [23] or *re-attentionizing* [26] and has not been explored much, even though it could be useful for various applications such as object enhancement [26, 24], directing viewers attention in mixed reality [25] or in computer games [3], distractor removal [13], background de-emphasis [28] and improving image aesthetics [29, 31, 15]. Imagine being able to highlight your child who stands in the chorus line, or making it easier for a person with a visual impairment to find an object by making it more salient. Such manipulations are the aim of this paper.

Professionals use complex manipulations to enhance a particular object in a photo. They combine effects such as increasing the object's exposure, decreasing the background




Figure 1: Our saliency driven image manipulation algorithm can increase or decrease the saliency of an image region. In this example the manipulation highlighted the bird while obscuring the leaf. This can be assessed both by viewing the image before (a) and after (c) manipulation, and by the corresponding saliency maps (b),(d).

exposure, changing hue, increasing saturation, or blurring the background. More importantly, they adapt the manipulation to each photo – if the object is too dark they increase its exposure, if its colors are too flat they increase its saturation etc. Such complex manipulations are difficult for novice users that often don't know what to change and how. Instead, we provide the non-experts an intuitive way to highlight (or obscure) objects. All they need to do is mark the target region and tune a single parameter, that is directly linked to the desired saliency contrast between the target region and the rest of the image. An example manipulation is presented in Figure 1.

Our approach has three key contributions when compared to previous approaches. First, previous methods focused mostly on highlighting a single image region, usually by re-coloring it with vibrant colors. The framework we propose is more generic. It handles multiple image regions and can either increase or decrease the saliency of each re-

gion. This generality is very useful, for example, to remove multiple background distractors, to highlight several important objects, or to change the focus from one object to another. Second, while most previous approaches are content with enhancing an object by recoloring it with a preeminent color, we aim to produce realistic and natural looking results. Obviously, coloring a region bright red in a gray-scale image would result in it being highly salient, however, the image would appear very non-realistic. We wish to avoid such results and manipulate the image in a way that is inline with its internal characteristics. Last, but not least, our approach provides the user with an intuitive way for controlling the level of enhancement/concealment. This important feature is completely missing from all previous methods.

The algorithm we propose aims at globally optimizing an overall objective that considers the image saliency map. A key component to our solution is replacing properties of image patches in the target regions with other patches from the same image. This concept is a key ingredient in many patch-bases synthesis and analysis methods, such as texture synthesis [11], image completion [1], highlighting irregularities [5], image summarization [27], image compositing and harmonization [9] and recently highlighting non-local variations [10]. Similar to these methods, we replace patches in the target regions with similar ones from other image regions. Differently from those methods, our patch-to-patch similarity considers the saliency of the patches with respect to the rest of the image. This is necessary to optimize the saliency-based objective we propose. A key observation we make is that these patch replacements do not merely copy the saliency of the source patch to the target location as saliency is a complex global phenomena. We therefore interleave saliency estimation within the patch synthesis process. In addition, we do not limit the editing to the target region but rather change (if necessary) the entire image to obtain the desired global saliency goal.

We assess our method by comparing two properties to previous methods: (i) The ability to manipulate an image such that the saliency map of the result matches the user goal. (ii) The realism of the manipulated image. These properties are evaluated via qualitative means, quantitative measures and user studies. Our experiments show a significant improvement achieved by our method.

2. Related Work

Previous work on attention retargeting had a mutual goal – to enhance a single selected region [24, 15, 28, 26, 29]. A good review and comparison of methods is provided in [23].

Briefly summarizing, the solutions proposed in [26, 24] focused mostly on color manipulation. In many cases they succeed to enhance the object of interest, but they also often produce non-realistic manipulations, such as purple snakes and blue flamingos. Other saliency cues, such as saturation,

illumination and sharpness are not well handled by these methods. [29, 15, 28, 25] do treat these cues but as our experiments show they often fail to achieve the desired manipulation.

Recently Yan et al. [31] suggested a deep convolutional network to learn transformations that adjust image aesthetics. One of the effects they study is Foreground Pop-Out, which is similar in spirit to object saliency enhancement. Their method produces aesthetic results, however, it requires intensive manual labeling by professional artists and is limited to the labeled effect (in their case, popup of a single foreground region).

Also related to our problem are methods that did not set their goal as saliency manipulation, however, their outcome effectively achieves this goal to some extent. Fried et al. [13] detect and remove distracting regions in an image via inpainting. Removing the distractors implicitly changes the image saliency map, however, it also alters the image composition. Instead, we attenuate the distractors so that they remain in the image but are not as salient. A somewhat related work [8] suggests a technique for camouflaging an object in a textured background by manipulating its texture. An aftereffect of their method is immersion of the object in the background, thus implicitly reducing its saliency. Their camouflage results are impressive but the approach is applicable mostly to certain types of textures.

3. Problem Formulation

Our image manipulation formulation generates an image J whose corresponding saliency map is denoted by S_J . It takes as input an image I , a target region mask R and the desired saliency contrast ΔS between the target region and the rest of the image. The user can also choose between three applications, as illustrated in Figure 2: (i) *Object Enhancement*, where the target is enhanced while the background saliency is decreased, (ii) *Distractor Attenuation*, where the target’s saliency is decreased, and (iii) *Background Decluttering*, where the target is unchanged while salient pixels in the background are demoted.

We pose this task as a patch-based optimization problem over the image J . The objective we define distinguishes between salient and non-salient patches and pushes for ma-




Figure 2: **Mask setups.** Illustration of the setups used for: (a) object enhancement, (b) getting rid of distractors and (c) decluttering. We increase the saliency in red, decrease it in blue and apply no change in gray.

nipulation that matches the saliency contrast ΔS . To do this we extract from I two databases of patches of size $w \times w$: $\mathcal{D}^+ = \{p; S_I(p) \geq \tau^+\}$ of patches p with *high* saliency and $\mathcal{D}^- = \{p; S_I(p) \leq \tau^-\}$ of patches p with *low* saliency. The thresholds τ^+ and τ^- are found via our optimization.

We start by formulating the problem of *Object Enhancement*. Later on we explain how this formulation is applied for the other setups. To increase the saliency of patches in R and decrease the saliency of patches outside R we define the following energy function:

$$\begin{aligned} E(J, \mathcal{D}^+, \mathcal{D}^-) &= E^+ + E^- + \lambda \cdot E^\nabla \quad (1) \\ E^+(J, \mathcal{D}^+) &= \sum_{q \in R} \min_{p \in \mathcal{D}^+} D(q, p) \\ E^-(J, \mathcal{D}^-) &= \sum_{q \notin R} \min_{p \in \mathcal{D}^-} D(q, p) \\ E^\nabla(J, I) &= \|\nabla J - \nabla I\|_2 \end{aligned}$$

where $D(p, q)$ is the sum of squared distances (SSD) over $\{L, a, b\}$ color channels between patches p and q . The first two terms suggest that patches in the target region R should be similar to patches in the subset \mathcal{D}^+ , i.e., have high saliency, while patches outside the target region R should have low saliency scores, as do the patches in the subset \mathcal{D}^- . The role of the third term, E^∇ , is to preserve the gradients of the original image I . The balance between the color channels and the gradient channels is controlled by λ .

The goal of optimizing the energy in (1) is to generate an image J with saliency map S_J , such that the contrast in saliency between R and the rest of the image is ΔS . The key to this lies in the construction of the databases \mathcal{D}^+ and \mathcal{D}^- . The higher the threshold τ^+ the more salient will be the patches in \mathcal{D}^+ and in return those in R . Similarly, the lower the threshold τ^- the less salient will be the patches in \mathcal{D}^- and in return those outside of R . Our algorithm performs an approximate greedy search over the thresholds to determine their values.

To formulate mathematically the affect of the user control parameter ΔS we further define a function $\psi(S_J, R)$ that computes the saliency difference between pixels in the target region R and those outside it:

$$\psi(S_J, R) = \text{mean}_{\text{top}\beta}\{S_J \in R\} - \text{mean}_{\text{top}\beta}\{S_J \notin R\} \quad (2)$$

and seek to minimize the saliency-based energy term:

$$E^{\text{sal}} = \|\psi(S_J, R) - \Delta S\| \quad (3)$$

For robustness to outliers we only consider the β ($= 20\%$) most salient pixels in R and outside R .

Adapting this formulation to other setups is trivial. In *Background Decluttering* we do not edit the target region by ignoring E^+ . For *Distractor Attenuation* we want to decrease the saliency of the target region so we simply invert the mask R and again ignore E^+ .

Algorithm 1 Saliency Manipulation

```

1: Input: Image  $I$ ; object mask  $R$ ; saliency contrast  $\Delta S$ .
2: Output: Manipulated image  $J$ .
3:
4: Initialize  $\tau^+$ ,  $\tau^-$  and  $J = I$ .
5: while  $\|\psi(S_J, R) - \Delta S\| > \epsilon$  do
6:   1. Database Update
7:      $\rightarrow$  Increase  $\tau^+$  and decrease  $\tau^-$ .
8:   2. Image Update
9:      $\rightarrow$  Minimize (1) w.r.t.  $J$ , holding  $\mathcal{D}^+, \mathcal{D}^-$  fixed.
10: end while
11: Fine-scale Refinement

```

4. Algorithm Overview

The optimization problem in (1) is non-convex with respect to the databases $\mathcal{D}^+, \mathcal{D}^-$. To solve it, we perform an approximate greedy search over the thresholds τ^+, τ^- to determine their values. Given a choice of threshold values, we construct the corresponding databases and then minimize the objective in (1) w.r.t. J , while keeping the databases fixed. Pseudo-code is provided in Algorithm 1.

Image Update: This step manipulates J according to the application setup selected by the user (Figure 2). Patches in regions to be enhanced are replaced with similar ones from \mathcal{D}^+ . Similarly, patches in regions to be demoted are replaced with similar ones from \mathcal{D}^- .

Database Update: This step reassigns the patches from the input image I into two databases, \mathcal{D}^+ and \mathcal{D}^- , of salient and non-salient patches, respectively. The databases are updated at every iteration by shifting the corresponding thresholds, in order to find values that yield the desired enhancement or concealment effects (according to ΔS).

Fine-scale Refinement: We observed that updating both J and $(\mathcal{D}^+, \mathcal{D}^-)$ at all scales does not contribute much to the results, as most changes happen already at coarse scales. Similar behavior was observed by [27] in retargeting and by [1] in reshuffling. Hence, the iterations of updating the image and databases are performed only at low-resolution. After convergence, we continue and apply the Image Update step at finer scales, while the databases are held fixed. Between scales, we down-sample the input image I to be of the same size as J , and then reassign the patches from the scaled I into \mathcal{D}^+ and \mathcal{D}^- using the current thresholds.

In our implementation we use a Gaussian pyramid with 0.5 scale gaps, and apply 5-20 iterations, more at coarse scales and less at fine scales. The coarsest scale is set to be 150 pixels width.

5. Detailed Description of the Algorithm

Throughout the algorithm when a saliency map is computed for either I or J we use a modified version of the

method of [21]. Because we want the saliency map to be as sharp as possible, we use a small patch size of 5×5 instead of the 9×9 in the original implementation. In addition, we omit the center prior which assumes higher saliency for patches at the center of the image. We found it to ambiguate the differences in saliency between patches, which might be good when comparing prediction results to smoothed ground-truth maps, but not for our purposes. We selected the saliency method of [21] since its core is to find what makes a patch distinct. It assigns a score $\in [0, 1]$ to each patch based on the inner statistics of the patches in the image, which is a beneficial property to our method.

Image Update In the Image Update step we minimize (1) with respect to J , while holding the databases fixed. This resembles the optimization proposed by [9] for patch-based image synthesis. It differs, however, in two important ways. First, unlike [9] that consider only luminance gradients, we consider gradients of all three $\{L, a, b\}$ color channels. This improves the smoothness of the color manipulation, preventing generation of spurious color edges, like those evident in Figure 3c. It guides the optimization to abide to the color gradients of the original image and often leads to improved results (Figure 3d).

As was shown in [9], the energy terms in (1) can be optimized by combining a patch *search-and-vote* scheme and a discrete Screened Poisson equation that was originally suggested by [4] for gradient domain problems. At each scale, every iteration starts with a *search-and-vote* scheme that replaces patches of color with similar ones from the appropriate patch database. For each patch $q \in J$ we *search* for the Nearest Neighbor patch p . Note that we perform two separate searches, for the target region and for the background, in either \mathcal{D}^+ (to increase saliency) or \mathcal{D}^- (to decrease saliency), depending on the application. This is the second difference from [9] where a single search is performed in one source region. To reduce computation time the databases are represented as two images: $I_{\mathcal{D}^+} = I \cap (S_I \geq \tau^+)$ and $I_{\mathcal{D}^-} = I \cap (S_I \leq \tau^-)$. The search is performed using PatchMatch [1] with patch size $w = 7 \times 7$ and translation transformation only (we found that rotation and scale were not needed for our application). In the *vote* step, every target pixel is assigned the mean color of all the patches that overlap with it. The voted color image is then combined with the original gradients of image I using a Screened Poisson solver to obtain the final colors of that iteration. We used $\lambda = 5$ as the gradients weight.

Having constructed a new image J , we compute its saliency map S_J to be used in the database update step explained next.

Database Update The purpose of the database update step is to search for the appropriate thresholds that split the patches of I into salient \mathcal{D}^+ and non-salient \mathcal{D}^- databases.

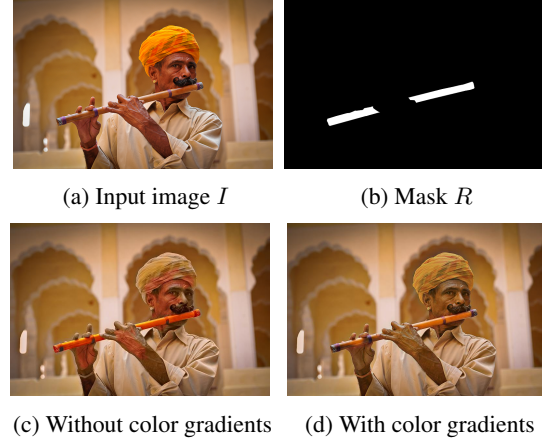


Figure 3: **Chromatic gradients.** A demonstration of the importance of chromatic gradients via an object enhancement example. (c) When *not* using color gradients - artifacts appear: orange regions on the flutist’s hat, hands and face. (d) By solving the screened Poisson equation on all three channels we improve the smoothness of the color manipulation, stopping it from generating spurious color edges, and the color of the flute is more natural.

Our underlying assumption is that there exist threshold values that result in minimizing the objective E^{sal} of (3).

Recall that the databases are constructed using two thresholds on the saliency map S_I such that $\mathcal{D}^+ = \{p; S_I(p) \geq \tau^+\}$ and $\mathcal{D}^- = \{p; S_I(p) \leq \tau^-\}$. An exhaustive search over all possible threshold values is non-tractable. Instead, we perform an approximate search that starts from a low value for τ^+ and a high value for τ^- and then gradually increase the first and reduce the second until satisfactory values are found. Note that \mathcal{D}^+ and \mathcal{D}^- could be overlapping if $\tau^+ < \tau^-$.

The naive thresholds $\tau^+ \approx 1$, $\tau^- \approx 0$, would leave only the most salient patches in \mathcal{D}^+ and the most non-salient in \mathcal{D}^- . This, however, could lead to non-realistic results and might not match the user’s input for a specific saliency contrast ΔS . To find a solution which considers realism and the user’s input we seek the maximal τ^- and minimal τ^+ that minimize the saliency term E^{sal} .

At each iteration we continue the search over the thresholds by gradually updating them:

$$\tau_{n+1}^+ = \tau_n^+ + \eta \cdot \|\psi(S_J, R) - \Delta S\| \quad (4)$$

$$\tau_{n+1}^- = \tau_n^- + \eta \cdot \|\psi(S_J, \bar{R}) - \Delta S\| \quad (5)$$

where \bar{R} is the inverse of the target region R . Since the values of the thresholds are not bounded, we trim them to be in the range of $[0, 1]$. Convergence is declared when $\|\psi - \Delta S\| < \epsilon$, i.e., when the desired contrast is reached. If convergence fails the iterations are stopped when the thresh-




Figure 4: Object Enhancement In all of these examples the user selected a single object region to be enhanced (b). To qualitatively assess the enhancement effect one should compare the input images in (a) to the manipulated images in (c,d,e), while considering the input mask in (b). The results of OHA in (c) are less realistic as they use arbitrary colors for enhancement. Since OHA is limited to hue changes in some cases it fails to enhance the object altogether (rows 7,8). WSR produces aesthetic results, but sometimes (e.g., rows 2,4 and 5) completely fails at enhancing the object. Our manipulation succeeds in enhancement while maintaining realism. In most images multiple enhancement effects occur simultaneously: emphasize by illumination (rows 2 and 8), emphasize by saturation (rows 2, 3 and 4) and emphasize by color (rows 1, 4-7).




Figure 5: **Enhancement evaluation:** The bars represent the (a) Correlation-Coefficient (CC) and (b) the Weighted F-beta (WFB) [22] scores obtained when comparing the ground-truth masks with saliency maps computed using five different saliency estimation algorithms (see text). The longer the bar, the more similar the saliency maps are to the ground-truth. It can be seen that the saliency maps of our manipulated images are consistently more similar to the ground-truth.

olds stop changing between subsequent iterations. In our implementation $\eta = 0.1$ and $\epsilon = 0.05$.

One important property of our method is that if $\tau^- = 1$ (or very high) and $\tau^+ = 0$ (or very low) we would get the image unchanged as the solution where all patches are replaced by themselves will lead to a zero error of our objective energy function (1).

Robustness to ΔS The only parameter we request the user to provide is ΔS . We argue that this parameter is easy to tune. For *Distractor Attenuation* and *Background Decluttering* we always set $\Delta S = 0.6$. For *Object Enhancement* the default value is $\Delta S = 0.6$, for which convergence was achieved in 90% of images. In only a few cases the result with $\Delta S = 0.6$ was not aesthetically pleasing and we used other values in the range $[0.4, 0.8]$. In the rest of the paper, if not mentioned otherwise, $\Delta S = 0.6$.

Our algorithm is not guaranteed to reach a global minima. However we found that typically the manipulated image is visually plausible, and pertains a good match to the desired saliency.

6. Empirical Evaluation

The assessment of our algorithm considers two properties of the manipulated image: its saliency map and whether it looks realistic. We compare our algorithm to OHA [24] and WSR [29], which stand out over other methods (according to [23]) in attention retargeting and aesthetics respectively¹. The comparison is both qualitative and quantitative.

¹Code for WSR is not publicly available, hence we used our own implementation. For OHA we use the original code.




Figure 6: **Realism evaluation.** Realism results obtained via a user survey (see text for details). Our manipulated images are ranked as more realistic than those of OHA and similar to those of WSR. The curves show the fraction of images with average score greater than *Realism score*. The Area-Under-Curve (AUC) values are presented in the legend.

The most time demanding part of our method is solving the screened Poisson equation at each iteration. Since our main focus was on quality we did not optimize the implementation for speed. Significant speed-up could be achieved by adopting the method of [12]. As was shown by [9] replacing these fast pyramidal convolutions with our current solver, will reduce run-time from minutes to several seconds.

6.1. Object Enhancement

We start by providing a qualitative sense of what our algorithm can achieve in Figure 4. Many more results are provided in the supplementary, and we encourage the reader to view them. Comparing to OHA, it is evident that our results are more realistic. OHA changes the hue of the selected object such that its new color is unique with respect to the color histogram of the rest of the image. This often results in unrealistic colors. The results of WSR, on the other hand, are typically realistic since their manipulation is limited in order to achieve aesthetic outcomes. This, however, comes at the expense of often failing to achieve the desired object enhancement.

The ability of our approach to simultaneously reduce and increase saliency of different regions is essential in some cases, e.g. Figure 4, rows 1 and 5. In addition, it is important to note that our manipulation latches onto the internal statistics of the image and emphasizes the objects via a combination of different saliency cues, such as color, saturation and illumination. Examples of these complex effects are presented in Figure 4, rows 3, 7 and 8, respectively.

Quantitative: To support these claims we further present quantitative evaluation on a corpus of 124 images gathered from [16, 2, 20, 13, 24, 19, 7] and from the web.

To measure how successful our manipulated images are, we do the following. We take the user provided mask as the ground-truth saliency map. We then compute saliency maps for each image using five different state-of-the-art methods:

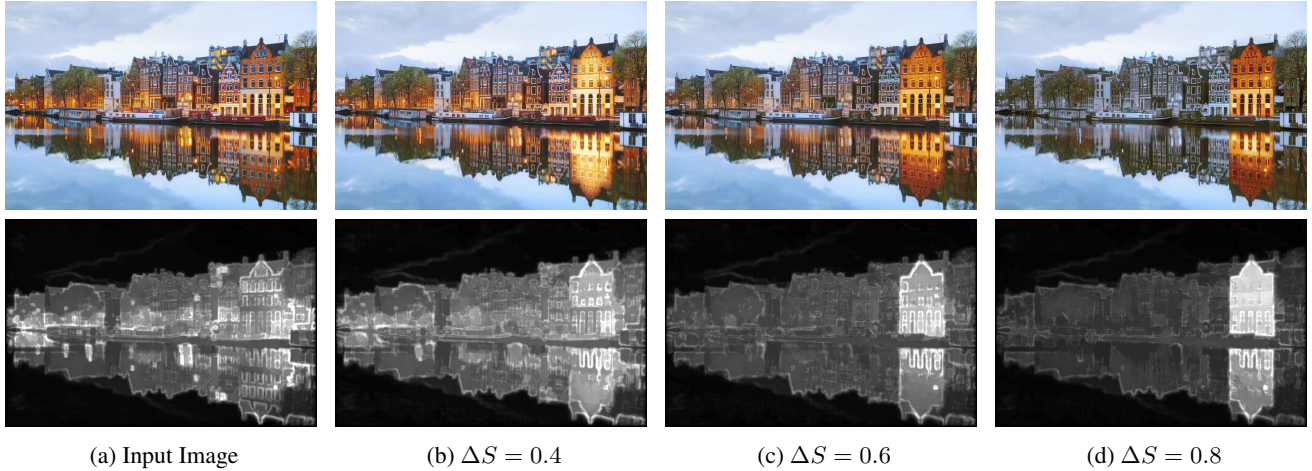


Figure 7: **Controlling the level of enhancement.** (a) Input image. The manipulated image J with $\Delta S = 0.4, 0.6, 0.8$, respectively. Bottom row shows the corresponding saliency maps. As ΔS is increased, so does the saliency contrast between the foreground and the background. As mask, the user marked the house and its reflection on the water.



Figure 8: **Background DeCluttering.** Often in cluttered scenes one would like to reduce the saliency of background regions to get a less noisy image. In such cases it suffices to loosely mark the foreground region as shown in (e), since the entire background is manipulated. In (a,b) saliency was reduced for the boxes on the left and red sari on the right. In (c,d) the signs in the background were demoted? to help draw attention to the bride and groom.



Figure 9: **Distractor Attenuation.** (a)-(c) Top: input images. The distractors were the balloon, the red flag, the shiny lamp and the red roof. Bottom: our manipulated images after reducing the saliency of the distractors. (e)-(f) Top: Zoom in on our result. Bottom: Zoom in on the inpainting result by Adobe Photoshop showing artifacts that are common to inpainting methods.

MBS [32], HSL [30], DSR [18], PCA [21] and MDP [17]. The computed saliency maps are compared to the ground-truth using two commonly-used metrics for saliency evaluation: (i) Pearson's-Correlation-Coefficient (CC) which was recommended by [6] as the best option for assessing saliency maps. (ii) Weighted F-beta (WFB) [22] which shown to be a preferred choice for evaluation of foreground maps. Figure 5 shows the saliency maps of our manipulated images are more similar to the ground-truth than those of OHA and WSR. This is true for both saliency measures and for all five methods for saliency estimation.

Realism: As mentioned earlier, being able to enhance a region is not enough. We would also like to verify that the manipulated images look plausible and realistic. We measure this via a user survey. Each image was presented to human participants who were asked a simple question: "Does the image look realistic?" The scores were given on a scale of [1-9], where 9 is 'definitely realistic' and 1 is 'definitely unrealistic'. We used Amazon Mechanical Turk to collect 20 annotations per image, where each worker viewed only one version of each image out of four. Figure 6 plots for each method the fraction of images with average score larger than a realism score ($\in [1, 9]$) and the overall AUC values. Our results are mostly realistic and similar to WSR, while OHA results are often non-realistic.

Controlling the Level of Enhancement: One of the advantages of our approach over previous ones is the control we provide the user over the degree of the manipulation effect. Our algorithm accepts a single parameter from the user, ΔS , which determines the level of enhancement. In Object Enhancement, the higher ΔS is, the more salient will the region of interest become. While we chose $\Delta S = 0.6$ for most images, another user could prefer other values to get more or less prominent effects. Figure 7 illustrates the influence ΔS on the manipulation results.

6.2. Other Applications

Distractor Attenuation: The task of getting rid of distractors was recently defined by Fried et al. [13]. Distractors are small localized regions that turned out salient against the photographer's intentions. In [13] distractors were removed entirely from the image and the holes were filled by inpainting. This approach has two main limitations. First, it completely removes objects from the image thus changing the scene in an obtrusive manner that might not be desired by the user. Second, hole-filling methods hallucinate data and sometimes produce weird effects.

Instead, we propose to keep the distractors in the image while reducing their saliency. Figure 9 presents some of our results and comparisons to those obtained by inpainting. We succeed to attenuate the saliency of the distractors, without having to remove them from the image.

Background Decluttering: Reducing saliency is also

useful for images of cluttered scenes where one's gaze dynamically shifts across the image to spurious salient locations. Some examples of this phenomena and how we attenuate it are presented in Figure 8. We refer to this task of reducing the clutter saliency as *saliency decluttering*.

This scenario is similar to that of removing distractors, with one main difference. Distractors are usually small localized objects, therefore, one could potentially use inpainting to completely remove them. Differently, when the background is cluttered, marking all the distractors could be tedious and removing them would result in a completely different image.

Our approach easily deals with cluttered background. The user is requested to loosely mark the foreground region. We then leave the foreground unchanged and manipulate only the background, using \mathcal{D}^- to automatically decrease the saliency of clutter pixels. The optimization modifies only background pixels with high saliency, since those with low saliency are represented in \mathcal{D}^- and therefore are matched to themselves.

7. Conclusions and Limitations

We propose a general visual saliency retargeting framework that handles multiple regions in the image. The method is able to manipulate an image to achieve a saliency change, while providing the user control over the level of change. The framework is applicable to various image editing tasks such as object enhancement, distractors attenuation and background decluttering. Comparison to previous work in terms of saliency manipulation and in terms of realism of the results shows a significant improvement.




Figure 10: **Limitations.** When the input image color set is narrow, the manipulation is limited, making it difficult to enhance a region.

Our method is not without limitations. First, since we rely on internal patch statistics, and do not augment the patch database with external images, the color transformations are limited to the color set of the image (see Figure 10). Second, since our method is not provided with semantic information, in some cases the manipulated image may be non-realistic. For example, in Figure 9, the balloon is colored in gray, which is an unlikely color in that context. Despite its limitations, our technique often produces visually appealing results that adhere to the user's wish.

We will release code for our algorithm.

References

- [1] C. Barnes, E. Shechtman, A. Finkelstein, and D. Goldman. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (TOG)*, 28(3):24, 2009. 2, 3, 4
- [2] S. Bell, K. Bala, and N. Snavely. Intrinsic images in the wild. *ACM Transactions on Graphics (TOG)*, 33(4):159, 2014. 6
- [3] M. Bernhard, L. Zhang, and M. Wimmer. Manipulating attention in computer games. In *Ivmsp workshop, 2011 ieee 10th*, pages 153–158. IEEE, 2011. 1
- [4] P. Bhat, B. Curless, M. Cohen, and C. L. Zitnick. Fourier analysis of the 2D screened poisson equation for gradient domain problems. In *ECCV 2008*, pages 114–128. Springer, 2008. 4
- [5] O. Boiman and M. Irani. Detecting irregularities in images and in video. *International Journal of Computer Vision*, 74(1):17–31, 2007. 2
- [6] Z. Bylinskii, T. Judd, A. Oliva, A. Torralba, and F. Durand. What do different evaluation metrics tell us about saliency models? *arXiv preprint arXiv:1604.03605*, 2016. 8
- [7] M.-M. Cheng, J. Warrell, W.-Y. Lin, S. Zheng, V. Vineet, and N. Crook. Efficient salient region detection with soft image abstraction. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1529–1536, 2013. 6
- [8] H.-K. Chu, W.-H. Hsu, N. J. Mitra, D. Cohen-Or, T.-T. Wong, and T.-Y. Lee. Camouflage images. *ACM Trans. Graph.*, 29(4):51–1, 2010. 2
- [9] S. Darabi, E. Shechtman, C. Barnes, D. B. Goldman, and P. Sen. Image Melding: Combining inconsistent images using patch-based synthesis. *ACM Transactions on Graphics (TOG)*, 31(4):82:1–82:10, 2012. 2, 4, 6
- [10] T. Dekel, T. Michaeli, M. Irani, and W. T. Freeman. Revealing and modifying non-local variations in a single image. *ACM Transactions on Graphics (TOG)*, 34(6):227, 2015. 2
- [11] A. A. Efros and T. K. Leung. Texture synthesis by non-parametric sampling. In *Proceedings of The IEEE International Conference on Computer Vision*, volume 2, pages 1033–1038. IEEE, 1999. 2
- [12] Z. Farbman, R. Fattal, and D. Lischinski. Convolution pyramids. *ACM Trans. Graph.*, 30(6):175, 2011. 6
- [13] O. Fried, E. Shechtman, D. B. Goldman, and A. Finkelstein. Finding distractors in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1703–1712, 2015. 1, 2, 6, 8
- [14] S. Goferman, L. Zelnik-Manor, and A. Tal. Context-aware saliency detection. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 34(10):1915–1926, 2012. 1
- [15] A. Hagiwara, A. Sugimoto, and K. Kawamoto. Saliency-based image editing for guiding visual attention. In *Proceedings of the 1st international workshop on pervasive eye tracking & mobile eye-based interaction*, pages 43–48. ACM, 2011. 1, 2
- [16] T. Judd, K. Ehinger, F. Durand, and A. Torralba. Learning to predict where humans look. In *IEEE International Conference on Computer Vision (ICCV)*, 2009. 6
- [17] G. Li and Y. Yu. Visual saliency based on multiscale deep features. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2015. 8
- [18] X. Li, H. Lu, L. Zhang, X. Ruan, and M.-H. Yang. Saliency detection via dense and sparse reconstruction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2976–2983, 2013. 1, 8
- [19] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014*, pages 740–755. Springer, 2014. 6
- [20] H. Liu and I. Heynderickx. Tud image quality database: Eye-tracking release 1, 2010. 6
- [21] R. Margolin, A. Tal, and L. Zelnik-Manor. What makes a patch distinct? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1139–1146, 2013. 1, 4, 8
- [22] R. Margolin, L. Zelnik-Manor, and A. Tal. How to evaluate foreground maps? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2014. 6, 8
- [23] V. A. Mateescu and I. Bajić. Visual attention retargeting. *IEEE MultiMedia*, 23(1):82–91, 2016. 1, 2, 6
- [24] V. A. Mateescu and I. V. Bajić. Attention retargeting by color manipulation in images. In *Proceedings of the 1st International Workshop on Perception Inspired Video Processing*, pages 15–20. ACM, 2014. 1, 2, 6
- [25] E. Mendez, S. Feiner, and D. Schmalstieg. Focus and context in mixed reality by modulating first order salient features. In *International Symposium on Smart Graphics*, pages 232–243. Springer, 2010. 1, 2
- [26] T. V. Nguyen, B. Ni, H. Liu, W. Xia, J. Luo, M. Kankanhalli, and S. Yan. Image Re-attentionizing. *Multimedia, IEEE Transactions on*, 15(8):1910–1919, 2013. 1, 2
- [27] D. Simakov, Y. Caspi, E. Shechtman, and M. Irani. Summarizing visual data using bidirectional similarity. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008. 2, 3
- [28] S. L. Su, F. Durand, and M. Agrawala. De-emphasis of distracting image regions using texture power maps. In *Proceedings of the 4th IEEE International Workshop on Texture Analysis and Synthesis*, pages 119–124. ACM, October 2005. 1, 2
- [29] L.-K. Wong and K.-L. Low. Saliency retargeting: An approach to enhance image aesthetics. In *IEEE Workshop on Applications of Computer Vision (WACV)*, pages 73–80. IEEE, 2011. 1, 2, 6
- [30] Q. Yan, L. Xu, J. Shi, and J. Jia. Hierarchical saliency detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1155–1162, 2013. 1, 8
- [31] Z. Yan, H. Zhang, B. Wang, S. Paris, and Y. Yu. Automatic photo adjustment using deep neural networks. *ACM Transactions on Graphics*, 2015. 1, 2
- [32] J. Zhang, S. Sclaroff, Z. Lin, X. Shen, B. Price, and R. Mech. Minimum barrier salient object detection at 80 FPS. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1404–1412, 2015. 1, 8